

From Autosubmit to Cylc

A community standard for the EC-Earth consortium?

Miguel Castrillo

Domingo Manubens

Lisbon, 23 October 2018

EC-Earth meeting

Running an EC-Earth experiment with Autosubmit



**Barcelona
Supercomputing
Center**

Centro Nacional de Supercomputación

Creating and configuring your experiment

```
> module load autosubmit
```

```
> module load ECAC
```

```
> autosubmit expid -H cca-intel -d "Test experiment to run EC-Earth with Autosubmit"
```

The new experiment "a1g4" has been registered.

Copying config files...

Experiment registered successfully

Remember to MODIFY the config files!

```
> cd /esarchive/autosubmit/a1g4/conf
```

Experiment configuration

Start dates, members and chunks (number and length).

Experiment project source: origin (version control system or path) and project configuration file path.

expdef_xxxx.conf

Workflow to be run: scripts to execute, dependencies between tasks, task requirements (processors, wallclock time...) and platform to use.

jobs_xxxx.conf

HPC, fat-nodes and supporting computers configuration.

Usually provided by technicians, users will only have to change login and accounting options for HPCs.

platforms_xxxx.conf

Project dependant experiment variables that Autosubmit will substitute in the scripts to be run.

proj_xxxx.conf



Editing up platform_XXXX.conf

[cca-intel]

Queue type. Options: ps, SGE, LSF, SLURM, PBS, eceaccess
eceaccess

TYPE = ecaccess

VERSION = pbs

HOST = cca

PROJECT = spesiccf

ADD_PROJECT_TO_HOST = false

USER = c3m

SCRATCH_DIR = /scratch/ms

TEST_SUITE = True

QUEUE = np

PROCESSORS_PER_NODE = 36

[ecmwf-cca]

Queue type. Options: ps, SGE, LSF, SLURM, PBS,

TYPE = ecaccess

VERSION = pbs

HOST = cca

PROJECT = spesiccf

ADD_PROJECT_TO_HOST = false

USER = c3a

SCRATCH_DIR = /scratch/ms

TEST_SUITE = True

QUEUE = np

PROCESSORS_PER_NODE = 36

Editing up expdef_XXXX.conf

[experiment]

Supply the list of start dates.

DATELIST =

Supply the list of members. Format fcX

MEMBERS =

Chunk size. NUMERIC = 4, 6, 12

CHUNKSIZE =

Total number of chunks in experiment. NUMERIC = 30, 15, 10

NUMCHUNKS =

[project]

If PROJECT_TYPE is set to none...

PROJECT_TYPE =

Destination folder name for project. type = STRING

PROJECT_DESTINATION =

[svn]

type = STRING, help = 'https://svn.ec-earth.org/ecearth3'

PROJECT_URL =

Select revision number. NUMERIC = 1778

PROJECT_REVISION =

[experiment]

Supply the list of start dates.

DATELIST = 19900101 19910101

Supply the list of members. Format fcX

MEMBERS = fc0 fc1

Chunk size. NUMERIC = 4, 6, 12

CHUNKSIZE = 1

Total number of chunks in experiment. NUMERIC = 30, 15, 10

NUMCHUNKS = 12

[project]

If PROJECT_TYPE is set to none...

PROJECT_TYPE = svn

Destination folder name for project. type = STRING

PROJECT_DESTINATION = ecearth3

[svn]

type = STRING, help = 'https://svn.ec-earth.org/ecearth3'

PROJECT_URL = https://svn.ec-earth.org/ecearth3/trunk

Select revision number. NUMERIC = 1778

PROJECT_REVISION = 6014

Editing up jobs_XXX.conf

[LOCAL_SETUP]

FILE = LOCAL_SETUP.sh

PLATFORM = LOCAL

[REMOTE_SETUP]

FILE = REMOTE_SETUP.sh

DEPENDENCIES = LOCAL_SETUP

WALLCLOCK = 00:05

[LOCAL_SETUP]

FILE = runtime/autosubmit/copy-runtime.sh

PLATFORM = LOCAL

[REMOTE_SETUP]

FILE = runtime/autosubmit/compilation.sh

DEPENDENCIES = LOCAL_SETUP

WALLCLOCK = 02:00

PROCESSORS = 18



Editing up jobs_XXX.conf

[SIM]

FILE = SIM.sh

DEPENDENCIES = INI SIM-1 CLEAN-2

RUNNING = chunk

WALLCLOCK = 00:05

PROCESSORS = 2

THREADS = 1

[TRANSFER]

FILE = TRANSFER.sh

PLATFORM = LOCAL

DEPENDENCIES = CLEAN

RUNNING = chunk

FREQUENCY = 5

[SIM]

FILE = runtime/autosubmit/ece-ifs+nemo.sh

DEPENDENCIES = REMOTE_SETUP SIM-1

RUNNING = chunk

WALLCLOCK = 01:00

PROCESSORS = 1:1:288:144

TASKS = 36

THREADS = 1

TASKS = 1:1:1:36

[TRANSFER]

FILE = runtime/autosubmit/transfer.sh

PLATFORM = LOCAL

DEPENDENCIES = SIM

RUNNING = chunk

Creating proj_XXX.conf

[common]

TEMPLATE_NAME = ecearth3

[grid]

IFS_resolution = T255L91

NEMO_resolution = ORCA1L75

[ioserver]

XIO_NUMPROC = 1

[ifs]

IFS_NUMPROC = 288

NFIXYR = 0

ATM_ini =

Creating proj_XXXX.conf

[nemo]

NEM_NUMPROC = 144

OCEAN_ini =

NEMO_remove_land = TRUE

[ice]

ICE = LIM3

ICE_ini =

Running your experiment

```
> autosubmit create a1g4
```

Preparing .lock file to avoid multiple instances with same expid.

Checking configuration files...

platforms_a1g4.conf OK

jobs_a1g4.conf OK

autosubmit_a1g4.conf OK

expdef_a1g4.conf OK

Configuration files OK

Using project folder: /esarchive/autosubmit/a1g4/proj

Loading parameters...

Creating the jobs list...

Creating jobs...

Adding dependencies...

Removing redundant dependencies...

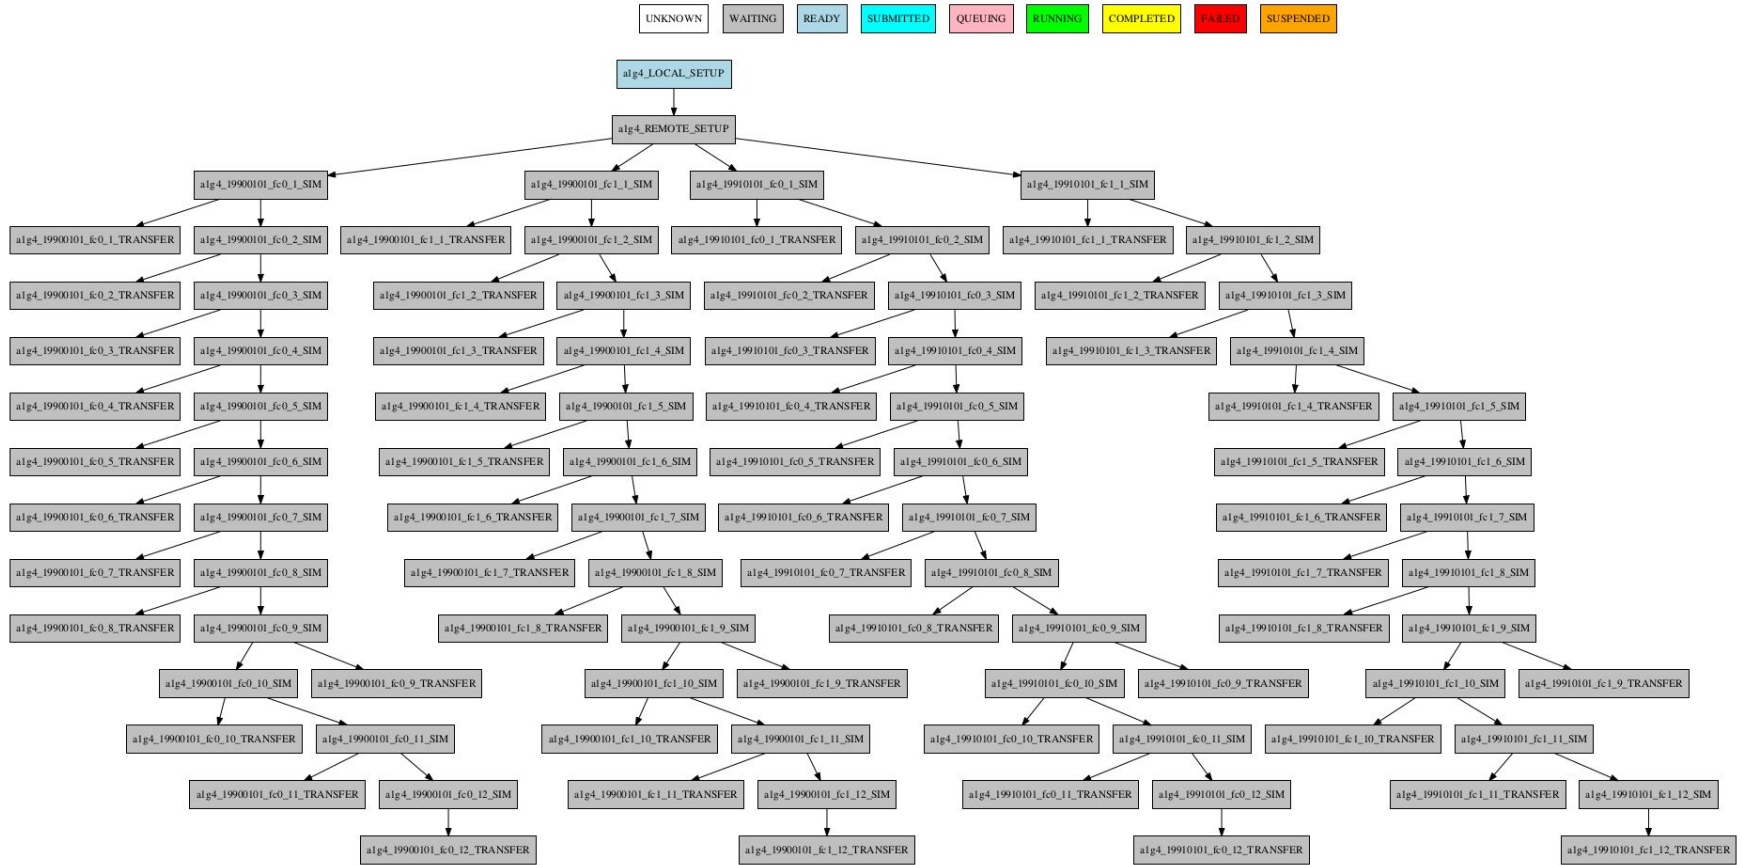
Saving the jobs list...

Plotting the jobs list...

Plotting...

```
> autosubmit run a1g4
```

... Sit back and relax ...



Changes in runtime/autosubmit

Index: copy-runtime.sh

=====

--- copy-runtime.sh (revision 6014)

```
../../sources/util/ec-conf/ec-conf --platform ${architecture} -o
PLT:ACTIVE:ECEARTH_SRC_DIR='${SCRATCH}/${exp_name}/auto-ecearth3/sources/sources' -o
PLT:ACTIVE:RUN_DIR='${SCRATCH}/${exp_name}/${run_start_date}/${member}/runtime' -o
PLT:ACTIVE:INI_DATA_DIR='${SCRATCH}/${exp_name}/${run_start_date}/${member}/inidata' config-run.xml
#../../sources/util/ec-conf/ec-conf --platform ${architecture} -o
PLT:ACTIVE:ECEARTH_SRC_DIR='${SCRATCH}/${exp_name}/auto-ecearth3/sources/sources' -o
PLT:ACTIVE:RUN_DIR='${SCRATCH}/${exp_name}/${run_start_date}/${member}/runtime' -o
PLT:ACTIVE:INI_DATA_DIR='/perm/ms/nl/nm6/ECE3-DATA' config-run.xml
```

Index: transfer.sh

=====

--- transfer.sh (revision 6014)

```
# local_transfer.sh: Copies the EC-Earth output from the
# remote machine to a local folder and extract it
```

```
LOCAL_DIR=/scratch/Earth/dmanuben/repository/exp/
```

Changes in runtime/autosubmit

Index: copy-runtime.sh

=====

+++ copy-runtime.sh (working copy)

```
#../sources/util/ec-conf/ec-conf --platform ${architecture} -o
PLT:ACTIVE:ECEARTH_SRC_DIR='${SCRATCH}/${exp_name}/auto-ecearth3/sources/sources' -o
PLT:ACTIVE:RUN_DIR='${SCRATCH}/${exp_name}/${run_start_date}/${member}/runtime' -o
PLT:ACTIVE:INI_DATA_DIR='${SCRATCH}/${exp_name}/${run_start_date}/${member}/inidata' config-run.xml
../sources/util/ec-conf/ec-conf --platform ${architecture} -o
PLT:ACTIVE:ECEARTH_SRC_DIR='${SCRATCH}/${exp_name}/auto-ecearth3/sources/sources' -o
PLT:ACTIVE:RUN_DIR='${SCRATCH}/${exp_name}/${run_start_date}/${member}/runtime' -o
PLT:ACTIVE:INI_DATA_DIR='/perm/ms/nl/nm6/ECE3-DATA' config-run.xml
```

Index: transfer.sh

=====

+++ transfer.sh (working copy)

```
# local_transfer.sh: Copies the EC-Earth output from the
# remote machine to a local folder and extract it
```

```
LOCAL_DIR=/esarchive/exp/ecearth/%EXPID%
```

All of this has been possible thanks to...

- **Convergence** between Autosubmit and classic **runtimes**
 - Can checkout any EC-Earth revision and run it with Autosubmit
- Recent changes in **ec-conf**
 - config-build.xml and config-run.xml can be changed parametrically
 - remove duplicate XML configuration files

From Autosubmit to Cylc



**Barcelona
Supercomputing
Center**

Centro Nacional de Supercomputación

EC-Earth use cases

We conducted a survey among the EC-Earth community to know about their use cases as also as their needs and expectations from a workflow management system.


34 participants

12 countries / 18 institutions

Country	Institution	# participants
Sweden	Lund (3), SU (2), SMHI (2)	7
Spain	BSC	6
Netherlands	KNMI (3), NSC (1), IMAU(1)	5
Finland	FMI (3), CSC (1)	4
Belgium	UCL	3
Germany	AWI (1), KIT (1)	2
Denmark	DMI	2
Italy	ENEA	1
Norway	UiB	1
Portugal	UL	1
United Kingdom	University of Oxford	1
Ireland	ICHEC	1


EC-Earth use cases

How do you run EC-Earth?




EC-Earth use cases

Model simulation production




EC-Earth use cases

Have you ever used a workflow management system?




How satisfactory your experience is?




- Very satisfied
- Satisfied
- Somewhat satisfied
- Not satisfied
- N/A

Would a workflow management system be useful to your work?




EC-Earth use cases

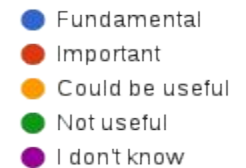

Database




E-mail notifications




Complex workflows




Jobs wrapper



GUI






Clock triggered jobs



Who is using Cylc?

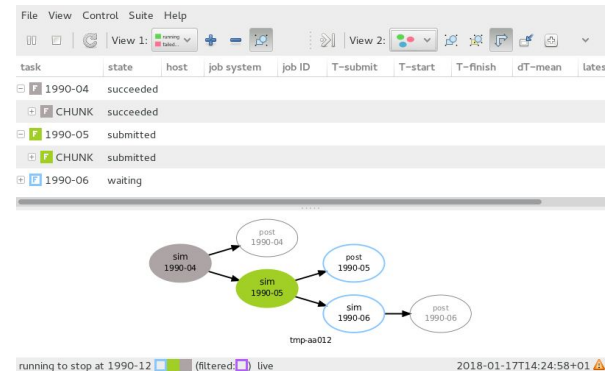
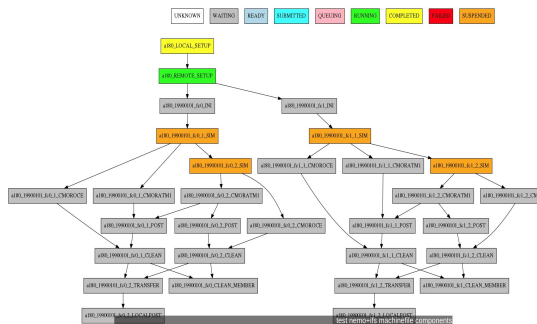
- National Centre for Water and Atmospheric Reserach (NIWA) - New Zealand
- Met Office - UK
- Max-Planck-Institut für Meteorologie - Germany
- Deutsches Klimarechenzentrum (DKRZ) - Germany
- Bureau of Meteorology - Australia
- NRL Marine Meteorology Division - USA
- Geophysical Fluid Dynamics Laboratory - USA
- Meteorological Service Singapore (MSS) - Singapore
- South African Weather Service (SAWS) - South Africa
- National Centre for Medium Range Weather Forecasting - India
- Korean Meteorological Administration (KMA) - Korea
- Centre of Excellence in Simulation of Weather and Climate in Europe - Europe
- National Center for Atmospheric Reserach (NCAR) - USA
- NOAA Environmental Modeling Center - USA
- Euro-Mediterranean Center on Climate Change (CMCC) - Italy
- Plymouth Marine Laboratory - UK
- Barcelona Supercomputing Center (BSC) - Spain
- National Centre for Atmospheric Science (NCAS) - UK
- Centre for Environmental Data Analysis (CEDA) - UK
- ARC Centre of Excellence for Climate Extremes (CLEX) - Australia

Cylc advantages over Autosubmit

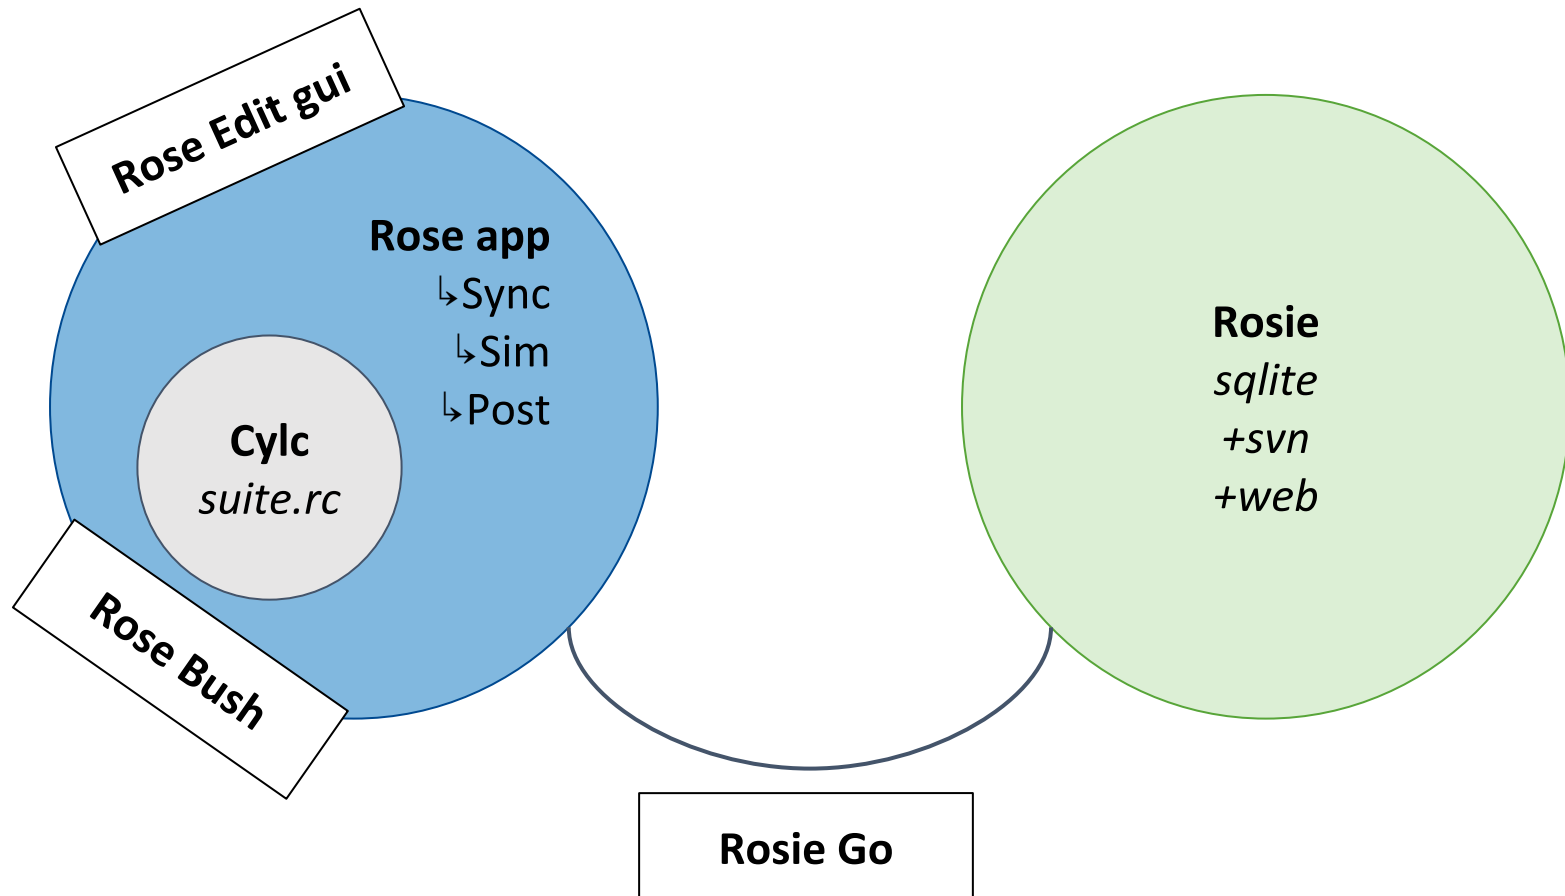
- Graphical user **interface** support → 80% 
- Better management for building **complex** workflows with families → 76% 
- Better **log** management
- Rose suite **configuration** tool + metadata for inputs
- Rose app: similar to proj_XXXX.conf and ec-conf but more powerful
- Rosie (svn): A frontend to **checkout** revisions of suites (identified uniquely), search and query suites and their metadata (Rose app + cylc suite) → 94% 
- Wider range of **users**
- IS-ENES3, ESIWACE **funding**
- MetO, NIWA, developers team **support**
- A lot of **documentation**: suite design guide, and tutorials
- **Available** on ECGATE and JASMIN

From Autosubmit to Cylc

- Option **Autosubmit** is **available** and will be **maintained**
- Option **Cylc** is being implemented



Cylc + Rose + Rosie



Rose - configuration management, GUI

Goal: to facilitate **suite development**, management and collaboration

The screenshot displays the Rose GUI interface. On the left is a sidebar with a tree view containing categories like 'suite info', 'suite conf', 'copy-runtime', 'env' (selected), 'post', 'sim', and 'transfer'. The main area is titled 'env' and shows 'Environment variable configuration' with fields for HPCARCH, HPCHOST, HPCROOTDIR, and HPCUSER. To the right, a panel titled 'Rose Bush @ bscearth200' shows user 'dmanuben' and suite 'tmp-aa012', with buttons for 'cycles list', 'task jobs list', 'broadcasts list', and 'cylc files'. Below this is a 'Display Options' section. At the bottom right, a table shows the suite's status across different cycle points.

Environment variable configuration

HPCARCH	marenostrum4
HPCHOST	mn1.bsc.es
HPCROOTDIR	/gpfs/scratch/bsc32/bsc32704/tmp-aa012
HPCUSER	bsc32704

Rose Bush @ bscearth200

dmanuben tmp-aa012

cycles list task jobs list broadcasts list cylc files

rose files rose-suite.info

Display Options

Suite ■ is stopped, last activity 6 minutes ago

cycle point	last active time	tasks	jobs	tasks	jobs	tasks	jobs	log/job-CYC
1990-04	6 minutes ago	0	0	✓ 1	⊙ 1	✗ 0	⊙ 0	
1990-03	11 minutes ago	0	0	✓ 2	⊙ 2	✗ 0	⊙ 0	
1990-02	17 minutes ago	0	0	✓ 2	⊙ 2	✗ 0	⊙ 0	
1990-01	22 minutes ago	0	0	✓ 4	⊙ 4	✗ 0	⊙ 0	

Rosie - experiment database

Goal: **version controlling suites** and suite **discovery** and management

tmp: Rosie Suites Discovery @ bscearth200

search query

Search term ...

☐ all revisions

search

detail	idx	branch	revision	owner	project	title	author	date	from_idx	status
	tmp-aa000	trunk	2	dmanuben	specs	test	eadmin	2017-10-18T08:29:22Z	None	A
	tmp-aa001	trunk	3	dmanuben	tmp	tmp1	eadmin	2017-11-27T12:25:58Z	None	A
	tmp-aa002	trunk	4	dmanuben	tmp	tmp	eadmin	2017-11-27T17:35:48Z	None	A
	tmp-aa003	trunk	5	dmanuben	tmp	tmp	eadmin	2017-11-27T17:52:51Z	None	A
	tmp-aa004	trunk	6	dmanuben	primavera	spinup	eadmin	2017-11-27T17:59:00Z	None	A
	tmp-aa005	trunk	7	dmanuben	specs2	test	eadmin	2017-11-27T18:14:03Z	tmp-aa000	A
	tmp-aa006	trunk	8	dmanuben	primavera2	spinup	eadmin	2017-11-28T09:39:04Z	tmp-aa004	A
	tmp-aa007	trunk	9	dmanuben	primavera3	spinup	eadmin	2017-11-28T09:39:39Z	tmp-aa006	A
	tmp-aa008	trunk	10	dmanuben	primavera4	spinup	eadmin	2017-11-28T09:39:57Z	tmp-aa007	A
	tmp-aa009	trunk	11	dmanuben	primavera5	spinup	eadmin	2017-11-28T09:41:09Z	tmp-aa008	A
	tmp-aa010	trunk	12	dmanuben	primavera6	spinup	eadmin	2017-11-28T09:41:41Z	tmp-aa009	A
	tmp-aa011	trunk	15	user1	blahblah	blah	user1	2017-11-29T10:15:03Z	None	M
	tmp-aa012	trunk	16	user1	EC-Earth	Test EC-Earth nemo only	user1	2017-11-29T14:35:54Z	tmp-aa011	A


Rose 2017.10.0

Current status

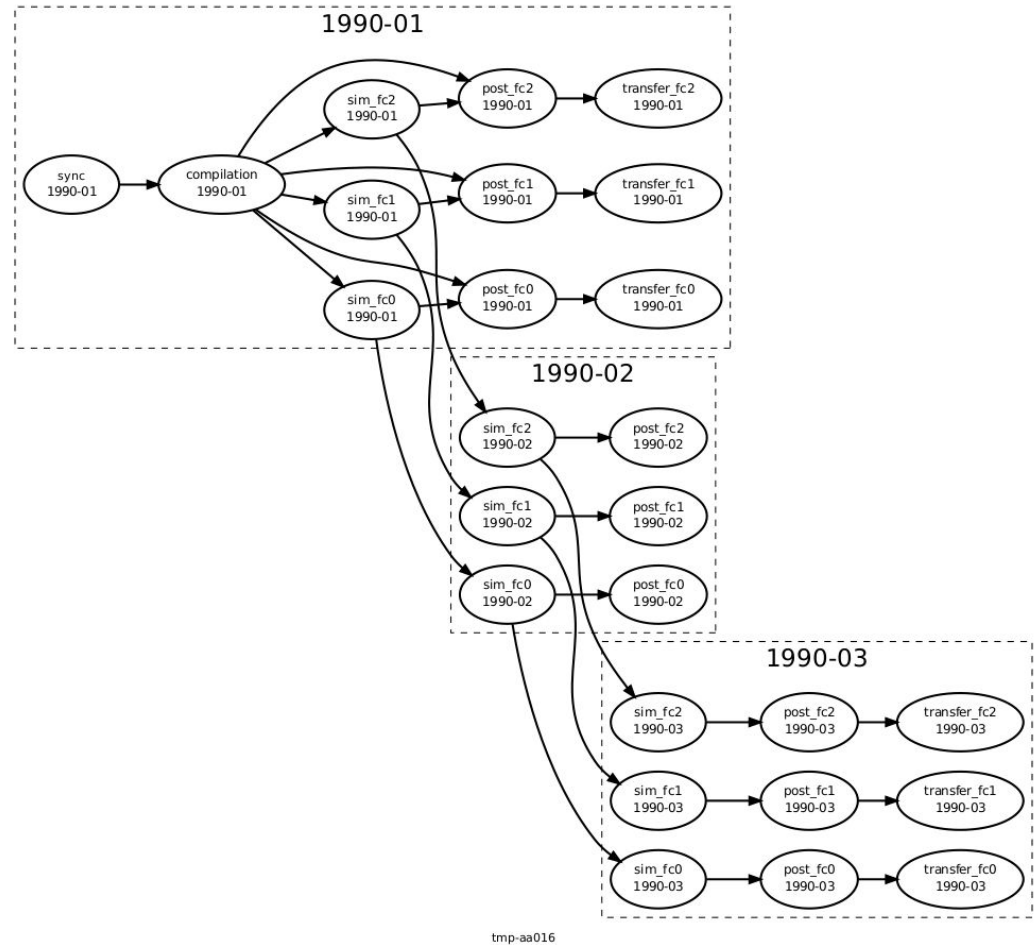

- @ BSC:
 - Cylc
 - Rose
 - Rosie: Not public
- @ ECMWF:
 - Cylc
 - Rose: Still under testing
 - Rosie: Need to configure an external server because the one at ecgate is for internal use

Cylc adoption at BSC (WIP)


- Climate prediction
 - ✓ • Historical runs
 - ✓ • Seasonal / decadal runs
 - Model development (stem?)
- Portable experiments including
 - ✓ • Compilation
 - Initialization
 - ✓ • Simulation
 - Post processing
 - Automated transfers
 - Data assimilation



Cylc workflow engine - Multi-member case



Ongoing work

- **Portable suites** for collaborative model development
 - EC-Earth **Cylc runtime** simplification
 - Rose **metadata** compliance
 - **Rosie web services** at EC-Earth Development portal? Alternatives?
- Cylc @ ECMWF
 - Rose suite **running at ECMWF**, including tasks to archive data at BSC and compute diagnostics on Power9 cluster
- (Next year) Single batch job for multiple tasks #2754 → 74% 

Summary

- A **strategy** for adopting Cylc and Rose as the EC-Earth common workflow management system is being build
 - Needs of the **community** have been taken into account
- A **demonstrator** has been created at BSC
 - Cylc, Rose and Rosie have been installed on a private virtual machine at BSC
 - Rose **suites** have been created for running EC-Earth **historical** and **multi-member** ensemble climate simulations
 - MareNostrum 4 (Slurm) and data-transfer nodes (background) are being used as remote resources, using polling

Next steps

- EC-Earth consortium **collaboration**
 - Develop more EC-Earth use cases
- Close **collaboration** with the MetOffice
 - ESiWACE 2
 - IS-ENES 3
- **Training** activities

Q & A



Thank you

“The research leading to these results has received funding from the EU H2020 Framework Programme under grant agreement n° H2020 GA 675191”

The content of this presentation reflects only the author’s view. The European Commission is not responsible for any use that may be made of the information it contains.

miguel.castrillo@bsc.es

domingo.manubens@bsc.es